GSMA | GTI

Mobile AI

Trust Paradigm Evolution for **Agentic Networks**

© GSMA 2026. All rights reserved.

Disclaimer

This document is not an Industry Specification and is not developed or maintained by GSMA in accordance with the provisions set out in GSMA AA.35 - Procedures for Industry Specifications.

This document does not represent the official position of the GSMA. The GSMA is the publisher and owner of this document, but views and opinions expressed in this report may be solely those of the author(s) and do not necessarily reflect the views of the GSMA.

This document is provided "AS IS", for information only. The GSMA makes no representation, warranty or undertaking (express or implied) with respect to and does not accept any responsibility for and disclaims any liability for the accuracy or completeness or timeliness of the information contained in this document. Any use of the information contained in this document is at the user's own risk and the user assumes liability for any third-party claims associated with such use. In no event shall the GSMA be liable for any direct, indirect, incidental, special, consequential or punitive damages, or any loss of profits, revenue, data or business opportunities arising out of or in connection with the use of or reliance on this document.

No licence is granted by implication, estoppel or otherwise under any intellectual property right of the GSMA or any third party. Nothing in this document shall be construed as granting any licence or right to use any trademark of the GSMA without prior written permission. Nothing in this document is intended to confer any benefit or right on any third party.

The GSMA is under no obligation to update, revise or supplement this document. The GSMA may withdraw or modify this document at any time without notice.

Antitrust Notice

The information contained herein has been prepared in accordance with the GSMA's Antitrust Policy: https://www.gsma.com/about-us/wp-content/uploads/2023/10/Antitrust_Policy_October_2023.pdf

Executive Summary

Driven by the rapid advancement of AI agents, next-generation intelligent networks are increasingly evolving from connecting people and things towards also connecting distributed intelligence. As network architectures and service paradigms continue to evolve, network environments may increasingly incorporate distributed, intelligent, and heterogeneous elements. While these developments unlock new capabilities, they also introduce new challenges related to security, trustworthiness, and governance. Zero Trust principles have been adopted in various conventional network environments, providing useful references for evolving network security architectures; however, as autonomous entities increasingly perform complex tasks, delegate actions, invoke network capabilities, and interact across different contexts and administrative domains, additional trust and authorisation considerations may arise beyond those addressed by conventional identity- and access-oriented implementations.

This white paper focuses on two categories of scenarios: first, scenarios in which external AI agents (Third-party AI Agents) invoke network capabilities exposed by MNOs; and second, AI for Networks scenarios within the MNO, in which Network AI Agents and AI-enabled network functions participate in network operation, management and orchestration.

While autonomous-network frameworks already support intent-driven operation and closed-loop automation, agentic interactions introduce additional requirements to ensure that delegated and evolving actions remain trusted and within the authorised objective. To address these challenges, this white paper proposes a new trust framework for Agentic Networks: the MAIC-Trust Framework (hereinafter referred to as M-Trust). The framework introduces three complementary dimensions for building a trustworthy agentic network environment:

Multi-party trust supports trust evaluation involving multiple relevant entities and trust-evidence sources, and allows different trust-establishment approaches to be applied according to the operational scenario. It provides a basis for collaborative trust evaluation and governance while preserving the security and policy authority of the relevant MNO and administrative domain.

Agentic Intent-aware trust focuses on the trustworthiness of intelligent agents themselves. It extends traditional identity verification mechanisms towards an integrated model of Identity Verification and Intent-aware Authorisation, enabling adaptive authorisation decisions that consider authorised objectives, execution context, resource requests, and risk conditions. This approach aims to strengthen security while avoiding unnecessary operational overhead.

Cross-domain trust addresses interactions across heterogeneous networks and administrative domains by supporting controlled exchange and evaluation of trust evidence under applicable operator policies and security requirements. This enables participating domains to establish and evaluate trust for cross-domain interactions while retaining authority over their own trust and authorisation decisions.

M-Trust is not intended to replace Zero Trust or existing mobile-network security mechanisms. It builds on relevant Zero Trust principles and considers how they may be complemented to address increasingly autonomous and agentic interactions in future mobile networks.

M-Trust is proposed as a reference framework and research direction for Agentic Networks rather than a standard specification or mandatory implementation architecture.

Contents

1 Evolution of Future Networks and Emerging Trust Requirements 02
1.1 From Connecting People and Things to Connecting Intelligence 02
1.2 Technology Drivers: New Architectures Create New Trust Demands 03
1.3 Business Drivers: Emerging Services Create New Trust Requirements 04
1.4 Trust Considerations for the Evolution towards Agentic Networks 05
1.5 From Traditional Trust to Zero Trust Architecture 05
1.6 Applicability and Limitations of Zero Trust Architecture 06
2 Core Concepts 08
2.1 Multi-party trust: Multi-Party Participation and Context-Aware Trust Establishment 09
2.2 Agentic Intent-aware trust: Intent-aware Authorisation and Adaptive Policy Management 10
2.3 Cross-domain trust: Cross-Domain Trust and Heterogeneous Trust Roots 11
2.4 Relationships with Zero Trust 13
3 High-Level Architecture 14
3.1 Security and Resilience of the M-Trust Framework 14
3.2 Overview of the Architecture 14
3.3 Trust Collection and Evaluation Layer: Data Foundation and Trust Perception 15
3.4 Policy Decision Layer: Policy Evaluation and Adaptation 17
3.5 Policy Enforcement Layer: Policy Enforcement and Response 18
3.6 Cross-Layer Collaboration and Trust Closed Loop 19
4 Technology Considerations for M-Trust 21
4.1 Relationship with Existing MNO Security Mechanisms 21
4.2 Protection and Assurance of Trust Evidence 22
4.3 Trust Evaluation and Policy Decision Techniques 23
4.4 Agent Identity, Delegation, and Secure Invocation 23
4.5 Cross-Domain Trust Evidence Exchange 24
4.6 Emerging and Use-Case-Specific Technologies 25
5 Conclusion and Future Outlook 26
5.1 Conclusion 26
5.2 Future Outlook 26
6 Abbreviations 28
7 Terms and Definitions 29
8 References 30

Foreword



Alex Sinclair
Chief Technology Officer
GSMA

In 1865, the United Kingdom passed the Locomotives Act, better known as the “Red Flag Act”. It required a person carrying a red flag to walk ahead of any self-propelled road vehicle to warn others of its approach. The intention was to manage the risks of a new technology, but the Act also became a lasting reminder that governance must evolve alongside innovation. Rules that do not keep pace hold back progress; innovation without effective safeguards undermines the confidence it needs to scale.

Artificial intelligence is reaching a similar point of transition. The risks of generative AI lie principally in what it says; the risks of agentic AI increasingly lie in what it does. As agents plan tasks, invoke tools and APIs, access business systems and delegate work across organisational and technical boundaries, trust can no longer be assessed only at the level of a single model or a single point in time.

In March 2026, the GSMA and GTI published the Mobile AI whitepaper. It set out a shared view of Mobile AI across four dimensions: AI for Networks, Networks for AI, Mobile AI terminals and agents, and Mobile AI applications, with security and trustworthiness as a core value. This whitepaper takes that commitment further by addressing an increasingly important question: how should trust evolve as networks begin to connect not only people and things, but also intelligent agents capable of taking action?

Established mechanisms for identity, authentication and access control remain essential. But it is not enough to know who or what an agent is. A trusted agentic environment must also determine who authorised an activity, what the agent is permitted to do, whether delegated actions remain consistent with the authorised objective, and whether permissions should be maintained, narrowed or withdrawn as circumstances change. This does not mean transferring authority from operators to AI. It means giving operators and other authorised parties better mechanisms to verify and govern how AI acts within their environments.

The MAIC-Trust framework presented here approaches this challenge through three complementary dimensions: multi-party trust, intent-aware trust and cross-domain trust. It proposes a closed loop in which trust and risk are reassessed as evidence, behaviour and context change, with authorisation remaining subject to operator policy, regulatory requirements and appropriate human oversight.

The mobile industry has a strong foundation for this work. The SIM established trusted identity and authentication at global scale. Roaming demonstrated how trust can extend across operator boundaries while each network retains control of its own domain. GSMA Open Gateway is creating a common, programmable way for developers to access network capabilities, enabling those capabilities to become part of increasingly intelligent and automated services. These are experiences the industry can build on as agents become part of networked services.

No organisation can build a trusted agentic environment alone. Agents will operate across networks, clouds, devices and applications. Open, interoperable approaches will be essential to avoid fragmentation while preserving each domain’s authority over its own trust decisions. That will require close collaboration among operators, technology companies, standards bodies, governments and the wider AI ecosystem.

Trust should not become a red flag held in front of innovation. Nor should innovation move ahead without effective safeguards. Our task is to create the frameworks, standards and shared practices that allow agentic technologies to move from isolated experiments to trusted, interoperable and scalable deployment.

I would like to thank all the organisations and experts who contributed to this whitepaper. I encourage the mobile and AI communities to build on its proposals, test them against real operational requirements, and work together to develop the trusted foundations for the next generation of intelligent networks.

1 Evolution of Future Networks and Emerging Trust Requirements

1.1 From Connecting People and Things to Connecting Intelligence

The evolution of communication networks has continuously expanded the boundaries of information interaction and system capabilities. Historically, three major stages can be identified:



Connecting People

In the early generations of mobile communications, networks primarily functioned as information transport infrastructures. Their mission was to enable voice, text, and multimedia communications between individuals, breaking geographical barriers and supporting human-to-human interaction.

Connecting Things

With the emergence of 5G and the Internet of Things (IoT), networks evolved towards connecting devices and physical environments. Sensors, industrial systems, and intelligent terminals became interconnected, enabling real-time mapping of physical entities into digital systems and supporting large-scale situational awareness.

Connecting Intelligence

As the industry moves towards the 6G and AI-Native era, the interaction between networks and AI may evolve along two complementary dimensions.The GSMA Mobile AI Whitepaper likewise identifies AI for Network and Network for AI as key dimensions of Mobile AI, and positions security and trustworthiness as one of the core values for its development [19].

First, Networks for AI enable distributed intelligence, including computing resources, AI models, and autonomous agents across cloud, edge, and endpoint environments, to communicate and collaborate. In this role, the network provides the connectivity and service capabilities required to support increasingly distributed AI workloads.

Second, AI for Networks refers to the increasing use of AI capabilities within the design, operation, management, and orchestration of networks themselves. In this context, intelligence becomes part of the capabilities used to operate and evolve the network, rather than solely an application or service running over it.

These two dimensions introduce distinct but related trust requirements. For Networks for AI, trust needs to extend beyond traditional identities and communication channels to interactions among distributed models, agents, data, and computing resources. For AI for Networks, trust also needs to address the use of AI in network operations and decision-making. The proposed framework therefore considers how trust mechanisms may evolve to support both AI-enabled services over the network and AI-enabled capabilities within the network itself.

Specifically, this white paper focuses on two categories of scenarios. The first is external AI agents invoking MNO network capabilities, for example a Third-party AI Agent requesting access to network capabilities exposed by an MNO through open interfaces such as GSMA Open Gateway. The second is AI for Networks within the MNO, in which Network AI Agents and AI-enabled network functions participate in network operation, management and orchestration. The former is a typical form of interaction between agents and MNO networks under the Networks for AI dimension, while the latter corresponds to the AI for Networks dimension. The trust requirements analysis and framework design in the following chapters are developed around these two categories of scenarios.

1.2 Technology Drivers: New Architectures Create New Trust Demands

Future communication networks are expected to undergo significant structural transformations in topology, operational intelligence, and deployment scope. Three relevant trends are increasingly distributed deployment, growing use of intelligence, and heterogeneous integration. Together, these developments challenge traditional security assumptions that rely on fixed boundaries and static identity verification.

Distributed Cloud and Edge Architectures

As computing capabilities become increasingly distributed across cloud and edge environments [1], the location and nature of security and trust boundaries may also evolve. However, the degree of decentralisation varies across different network layers and deployment scenarios, and clearly defined boundaries remain important in areas such as network interconnection. Trust mechanisms therefore need to accommodate both established boundary-based security controls and more distributed, continuously adaptive approaches, depending on the architecture and use case.

Intelligent Architectures

AI agents are expected to become integral components of both network functions and user devices [2]. Network entities will evolve from purely passive execution elements into more intelligent entities with varying degrees of perception, reasoning, collaboration, and autonomous decision-making capabilities. In this environment, identity verification remains necessary but may need to be complemented by evaluation of agent behaviour, authorisation decisions, execution contexts, computational integrity, and multi-agent collaboration.

Heterogeneous Integrated Networks

Future integrated terrestrial, aerial, satellite, and maritime networks will combine diverse communication systems, protocols, governance models, and trust domains. Such environments introduce significant complexity due to heterogeneous security policies,

dynamic connectivity conditions, and cross-domain interactions. Trust architectures may therefore need to support the controlled exchange and evaluation of trust evidence across permitted administrative boundaries, while maintaining continuity and resilience under varying operational conditions.

In summary, architectural evolution not only improves network capabilities but also exposes limitations in existing security frameworks. Existing trust and security mechanisms may need to be complemented to address increasingly autonomous, distributed, and cross-domain interactions. Future trust frameworks should support these interactions while preserving established MNO security controls, administrative authority, and regulatory boundaries.

1.3 Business Drivers: Emerging Services Create New Trust Requirements

Beyond architectural evolution, emerging service scenarios are also reshaping trust requirements. The ITU has defined six usage scenarios for IMT-2030 [3], three of which extend the existing IMT-2020 scenarios while the remaining three are new, forward-looking scenarios. Among these, three particularly important future service domains are expected to drive new trust demands:



Artificial Intelligence and Communication

As AI becomes deeply integrated into communication systems, trust relationships will increasingly extend beyond humans and devices to encompass AI-enabled network functions, services and autonomous agents, including interactions among autonomous agents and between humans and agents. Because a growing share of decisions is carried out autonomously by agents, systems need to support identity recognition and intent alignment among heterogeneous agents without human intervention in every interaction. Given the potential opacity of AI models, appropriate mechanisms are needed to enable critical autonomous decisions to be traced, assessed, and, where necessary, reviewed by relevant stakeholders.



Integrated Sensing and Communication

Integrated sensing and communication systems enable networks not only to transmit information but also to observe, interpret, and interact with physical environments [4]. While this significantly expands network capabilities, it also introduces additional privacy and security considerations that need to be appropriately addressed. Future trust frameworks must provide safeguards throughout the entire lifecycle of sensing data, including collection, storage, processing, sharing, and destruction. Sensing and communication systems should support appropriate authorisation, purpose-limitation, access-control, and data-governance mechanisms that enable data to be used securely across multiple stakeholders while protecting sensitive information [15].



Ubiquitous Connectivity

Future communication environments will need to support seamless interactions across terrestrial, satellite, aerial, maritime, and private networks. Users, devices, and agents may move across organizational, technological, and geographical boundaries. These scenarios may require mechanisms through which relevant credentials and trust evidence can be securely conveyed and evaluated across permitted domain boundaries. Evidence established in one domain may inform another domain's decision, but acceptance remains subject to the receiving domain's policies, security controls, and regulatory obligations. MNO-anchored identity, built on SIM and eSIM credentials and network-based authentication, provides a trusted foundation for trust mobility and cross-domain collaboration. This foundation can be extended through standardised interfaces such as GSMA Open Gateway APIs, and complemented where appropriate by federated identity frameworks.

1.4 Trust Considerations for the Evolution towards Agentic Networks

The evolution towards increasingly autonomous networks already introduces intent-driven operation, AI-assisted decision-making, and closed-loop automation to improve network and service management. M-Trust is not intended to replace or duplicate these capabilities. Instead, it addresses an additional trust and authorisation challenge that becomes increasingly important as autonomous network functions evolve towards agentic operation.

In an Agentic Network, an authorised objective may be fulfilled through multiple autonomous actions, including task decomposition, delegation to other agents, invocation of network capabilities, and interactions across changing execution contexts. The security question therefore extends beyond whether an autonomous system can achieve the intended operational outcome. MNOs also need to determine whether the entities involved remain trusted and whether the actions selected during execution remain within the scope of the authorised objective, applicable operator policy, and granted authorisation.

M-Trust therefore complements autonomous network operations by supporting continuous evaluation of whether agent actions remain trusted and authorised as objectives are decomposed, delegated, and executed. Changes in trust, risk, or context may lead to bounded adjustment, additional verification, or operator-controlled intervention within existing MNO security and administrative authority.

1.5 From Traditional Trust to Zero Trust Architecture

Conventional network security has relied largely on perimeter-based trust models. Such models divide networks into security domains and use mechanisms such as firewalls, virtual private networks, network segmentation and access control to protect and restrict access across security boundaries. In these models, network location and boundaries typically carry significant trust weight, and entities that have entered a trusted network domain may be granted a degree of implicit trust. In mobile networks, trust relationships have likewise long been built on well-defined identities, credentials, pre-configured security relationships and operator administrative boundaries. Examples include user and device trust based on subscription relationships, identity credentials and access authentication mechanisms; cross-domain trust between operators established through roaming and interconnection relationships; and trust among network functions established on the basis of defined security relationships and operator policy. In network architectures that are relatively closed and have clearly defined administrative boundaries, these mechanisms have provided an effective security foundation for mobile networks. However, as networks become increasingly cloud-based and service-based, and as capability exposure and multi-party collaboration expand, traditional network boundaries are becoming more complex, and relying solely on network location or pre-established trust relationships is no longer sufficient to address evolving security risks. Once an attacker breaches the perimeter or obtains legitimate identity credentials, they may attempt to access other internal resources or move laterally. In addition, traditional security models tend to rely more on authentication and authorisation at initial access or connection establishment, with relatively limited consideration of how subject behaviour, execution context and risk status continue to change

after a session is established. Zero Trust has developed against this background as a security paradigm.

The concept of Zero Trust was systematically introduced by John Kindervag, an analyst at Forrester Research, in 2010. In 2020, the National Institute of Standards and Technology (NIST) published SP 800-207, Zero Trust Architecture [5], which defined an abstract Zero Trust Architecture and described its logical components and deployment models. Based on NIST SP 800-207 and subsequent research, Zero Trust can be summarised through three fundamental principles:


Draft NIST Special Publication 800-207

Zero Trust Architecture

Scott Rose
Oliver Borchert
Stu Mitchell
Sean Connelly

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-207-draft

COMPUTER SECURITY

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

Identity-Centric Security

Zero Trust places identity at the centre of security control. Access decisions require appropriate authentication and authorisation of the relevant subject and device or workload, together with applicable contextual information.

Within MNO networks, the same principle applies to the network itself, not only to the entities that use it. In the 5G Service-Based Architecture, network-function communication is protected through established authentication, authorisation and interface-security mechanisms, while subscriber credentials and access authentication are protected through the relevant 3GPP procedures. Extending identity-centric controls to inter-operator interconnection, network management and operations access, and network equipment assurance allows operators to apply Zero Trust consistently across both the services they deliver and the infrastructure that delivers them.

Least Privilege

The principle of least privilege ensures that entities are granted only the minimum permissions necessary to accomplish a specific task. By limiting the scope of access rights, Zero Trust reduces the likelihood of lateral movement and minimises the potential impact of security breaches. Risk and context may support the restriction of permissions or the application of additional conditions. Any relaxation of controls should remain within explicitly defined policy and authorisation boundaries.

Continuous Monitoring

Zero Trust continuously monitors system behaviour, traffic patterns, and operational status to identify anomalies and emerging threats. Through real-time monitoring and adaptive access control, the architecture enables continuous assessment and mitigation of security risks, supporting timely security response and operational resilience.

Through these mechanisms, Zero Trust shifts security from traditional perimeter-based protection to continuous protection driven by identity and behavioural verification. Its principles have been applied in range of cloud-native architectures, virtualised environments, and multi-tenant platforms.

1.6 Applicability and Limitations of Zero Trust Architecture

The applicability of Zero Trust to telecommunications networks has been explored extensively. For example, 3GPP TR 33.894 studies the applicability of selected Zero Trust security principles to mobile networks [6], while TR 33.794 studies enablers including data exposure for security evaluation and monitoring and mechanisms for dynamic policy enforcement [7]. These studies demonstrate the relevance of Zero Trust principles to strengthening security in telecommunications networks. However, the work has primarily remained at the study and enabling-mechanism level, and Zero Trust has not been established as a comprehensive normative security architecture for the 5G System.

Despite these benefits, existing Zero Trust architectures are primarily designed for users, devices, workloads, and resource access. As networks evolve towards agent-centric environments characterised by autonomous decision-making, dynamic cooperation, tool invocation, delegated permissions, persistent memory, probabilistic reasoning, and cross-domain interactions, several limitations become apparent.

Scalability and Efficiency of Continuous Trust Evaluation

Zero Trust supports continuous monitoring and policy evaluation, together with re-authentication or re-authorisation where required by policy and context based on the principle of least privilege. These principles remain important in AI agent environments, where autonomous agents may dynamically interact with multiple services, tools, and other agents. However, applying conventional verification and policy evaluation mechanisms uniformly to every high-frequency agent interaction may introduce computational and communication overhead and may not scale efficiently in latency-sensitive scenarios. The challenge is therefore not to relax least-privilege or verification requirements, but to make trust evaluation more adaptive, contextual, and risk-aware, enabling appropriate security controls to be applied dynamically according to the identity, the authorised objective, execution context, requested network capability, and associated risk.

Applying Deterministic Security Controls to Probabilistic Agent Behaviour

AI agents rely on probabilistic reasoning influenced by model characteristics, available data, task context, and interactions with external tools and environments. Their specific decisions and execution paths therefore cannot always be predicted in advance. This uncertainty does not reduce the need for deterministic security controls; rather, it reinforces the importance of clearly defined authorisation boundaries, least-privilege access, and enforceable constraints on agent actions. However, static rules based only on identity and predefined access conditions may lack sufficient context to assess dynamically generated actions and changing risk conditions. The challenge is therefore to combine deterministic policy enforcement with continuous, context-, intent-, and risk-aware trust evaluation. Under this approach, adaptive trust signals may inform whether access is granted, constrained, escalated for additional approval, or denied, while the resulting security decision remains enforceable, auditable, and subject to predefined governance boundaries.

Limited Adaptability Across Mobile Network Trust Domains

Future mobile networks are expected to support increasingly autonomous AI agents operating across network functions, edge and cloud environments, service platforms, and multiple administrative domains. For MNOs, such interactions must operate within clearly defined security, policy, regulatory, and administrative boundaries. Conventional Zero Trust approaches are typically applied within a defined organizational or trust domain, with policy decisions governed by the responsible domain authority. As agentic capabilities become more distributed and interact across operator, enterprise, cloud, and service-provider environments, establishing and evaluating trust across these boundaries becomes increasingly complex. Future mobile networks therefore require mechanisms that can establish and convey verifiable trust context across permitted domain boundaries, while preserving the policy authority, regulatory obligations, and security controls of each participating domain.

2 Core Concepts

Traditional Zero Trust architectures are built upon three foundational principles: Identity-Centric Security, Least Privilege, and Continuous Monitoring. These principles have provided a robust foundation for modern network security [8]. However, the introduction of agentic capabilities into future telecommunications networks may enable greater autonomy in network operations, more dynamic service interactions, distributed collaboration, and cross-domain coordination. Given the critical-infrastructure role of telecommunications networks, such autonomy must operate within clearly defined security, policy, and governance boundaries. GSMA security work has also identified Generative and Agentic AI as an emerging area requiring continued security attention in the mobile ecosystem [16]. This creates additional trust requirements, including the ability to continuously assess agent identity, intent, context, and risk while preserving deterministic authorisation boundaries, least-privilege access, accountability, and operator control.



Figure 1: M-Trust extends Zero Trust

As illustrated in Figure 1, MAIC-Trust (hereinafter referred to as M-Trust) extends Zero Trust and introduces three additional dimensions:

Multi-party trust: supporting relevant multi-party evidence and context-appropriate trust establishment.

Agentic Intent-aware trust: supporting intent-aware authorisation and bounded policy adaptation.

Cross-domain trust: supporting trust anchors and evidence, and controlled evidence exchange across permitted administrative boundaries.

Working together, these pillars provide a scalable, adaptive, and collaborative trust foundation for next-generation intelligent communication networks.

2.1 Multi-party trust: Multi-Party Participation and Context-Aware Trust Establishment

Multi-party trust addresses scenarios in which relevant trust evidence is held by multiple authorised parties or where no single participant controls the complete trust relationship. It supports multi-party evidence contribution and context-appropriate trust-establishment approaches while preserving each MNO's policy and authorisation authority.

In this white paper, Multi-party trust refers to the combination of federated multi-party trust evidence and evaluation, and context-aware trust establishment [9], providing a flexible and reliable foundation for distributed collaboration.

Federated Multi-Party Trust Evidence and Evaluation

In multi-operator and cross-domain telecommunications environments, relevant trust evidence may be distributed across MNOs, roaming and interconnect partners, edge service providers, application providers, enterprise customers, and other authorised participants. For example, a home MNO may hold identity and subscription-related evidence, a visited MNO may observe local network behaviour, an edge provider may provide workload or platform attestation, and an application provider may provide information about an agent's software integrity, delegated authority, and requested task.

The purpose of multi-party trust evaluation is not to require all participating entities to reach the same trust judgement or to replace a single trust authority with a single globally aggregated trust outcome. Instead, authorised parties may contribute independently verifiable, purpose-limited and time-bounded trust evidence. Each MNO retains control of its own trust and authorisation decisions and evaluates the available evidence according to its service context, security policies, regulatory obligations, risk appetite, and relationship with the requesting party.

External evidence should therefore be treated as an input to, rather than a substitute for, the MNO's local policy decision. Evidence may be validated and weighted according to its provenance, integrity, freshness, relevance, independence, and the assurance level of its source. Depending on the risk of the requested action, it may instead trigger restricted access, additional verification, human or operator approval, isolation, or denial. However, adding more evidence providers also introduces additional risks, including compromised evaluators, coordinated or collusive manipulation of trust evidence, correlated evidence, stale information, and manipulation of aggregation mechanisms. Implementations should protect the authenticity, integrity, freshness and independence of contributed evidence, support revocation and auditability, and retain fail-safe local policy enforcement. Where trust evidence is exchanged between independent market participants, appropriate safeguards should also prevent the exchange of competitively sensitive information or the use of shared trust mechanisms to unfairly exclude third parties.

Consensus mechanisms may be applicable in cases where participants, fault assumptions, decision scope, and governance rules are explicitly established. They should not be assumed to produce a universal trust outcome across all MNOs or services.

Context-Aware Trust Establishment

Different interactions in future mobile networks may require different approaches to trust establishment. For example, interactions between AI agents and network functions within a

single MNO administrative domain may rely primarily on operator-controlled identity, policy and authorisation mechanisms. When a third-party AI agent requests access to exposed network capabilities, additional trust evidence or endorsement from an authorised external party may be required, subject to the MNO's security, policy and regulatory requirements. In multi-operator or other federated scenarios, where no single operator controls the entire trust relationship, federated or distributed trust mechanisms may be needed to establish verifiable trust among participating domains. Context-Aware Trust Establishment therefore does not imply dynamically replacing the underlying MNO security architecture; rather, it enables the appropriate trust-establishment mechanism to be applied according to the entities involved, administrative boundaries, risk level, and applicable policy and regulatory requirements.

2.2 Agentic Intent-aware trust: Intent-aware Authorisation and Adaptive Policy Management

Agentic Intent-aware trust is designed for next-generation Agentic Network architectures. It adopts Identity Verification and Intent-aware Authorisation as its foundational mechanisms. By jointly considering task objectives, execution context, resource requests, and risk conditions, Agentic Intent-aware trust enables dynamic security decision-making. In this framework, intent refers to an expressed objective or task expectation, including relevant goals and constraints. The actions derived by an agent in pursuit of that intent remain separately subject to authorisation and policy enforcement. The framework addresses the autonomous and dynamic interaction characteristics of intelligent agents through two complementary capabilities. Rather than assuming that an agent's future actions can be fully predicted, the framework continuously evaluates observable trust signals, declared intent, execution context, and risk conditions to inform security decisions. Agent actions remain subject to predefined, enforceable security and policy boundaries, with additional verification, privilege restriction, escalation, or denial applied where trust evidence is insufficient or risk exceeds acceptable thresholds.

Intent-aware Authorisation

Existing access-control mechanisms can support dynamic and policy-based decisions. Agentic Intent-aware trust considers additional agent-specific context, including the authorised objective, delegation and invocation chains, execution context, and observed behaviour.

In multi-agent environments, agent behaviour may be dynamic as agents adapt their plans and propose different actions in response to changing task conditions. However, such adaptation must not imply autonomous expansion of an agent's authorisation scope. In telecommunications networks, particularly where actions may affect sensitive data or critical network functions, any proposed action that exceeds the agent's existing authorisation scope should be subject to additional policy evaluation and, where required, explicit approval before execution. Consequently, authorisation decisions should continuously evaluate the alignment between the authorised objectives, requested actions, execution context, and actual resource usage, while dynamically adapting authorisation according to changing trust and risk conditions. An agent may adapt its execution plan, and its effective permissions may be automatically adjusted within predefined policy, task, and risk boundaries. However, the agent itself should not be able to autonomously expand its privileges beyond those boundaries.

The framework continuously evaluates task context, invocation chains, resource requests, observable behavioural characteristics, and risk signals to determine whether a requested action remains consistent with the authorised objective and applicable security policies. For low-risk actions that fall within a pre-authorised policy envelope, the system may automatically grant the minimum, time- and scope-bounded permissions required for execution. As trust or risk conditions change, those permissions may also be restricted, revoked, or subject to additional verification.

For actions outside the predefined authorisation envelope, particularly cross-domain access to sensitive data, core network configuration changes, or other high-impact operations, a new authorisation decision should be triggered. Depending on the applicable MNO policy and risk level, this may require stronger assurance, explicit approval, additional constraints, or denial.

Automated and Adaptive Policy Management

Automated and Adaptive Policy Management focuses on automating and dynamically adapting trust-policy enforcement within operator-defined governance, security, and regulatory boundaries. Rather than making the entire policy lifecycle fully autonomous, it enables policies and associated controls to be selected, applied, and adjusted according to changing trust, risk, and operational conditions, while policy authority and governance remain under MNO control.

Unlike traditional policy management approaches that depend heavily on manual configuration and periodic updates, Agentic Intent-aware trust continuously collects environmental information, including risk conditions, data sensitivity levels, and task urgency, together with agent behaviour information such as compliance history, task completion records, task performance, historical behavioural consistency, anomalous activity frequency, and network resource usage patterns. These inputs are aggregated into dynamic trust and risk assessments. Based on these assessments, together with the authorised objective, execution context, network capability and resource sensitivity, and operator-defined policy constraints, the system may automatically select, apply, or adapt access-control conditions and risk-mitigation measures within predefined policy boundaries. Lower-risk operations may use streamlined verification procedures where the applicable baseline security and authorisation requirements remain satisfied, whereas high-risk actions involving sensitive network capabilities and resources or cross-domain interactions may trigger enhanced verification, fine-grained authorisation, and continuous auditing or additional operator-controlled approval where required.

Policies are not static. As environmental conditions evolve, risks increase, or abnormal behaviour emerges, the system can automatically strengthen monitoring, restrict authorisation scopes, and trigger additional safeguards within operator-defined constraints. Through this dynamic synchronisation between risk conditions and policy enforcement, Agentic Intent-aware trust can improve responsiveness to changing operating conditions while preserving MNO policy authority, security controls, and applicable regulatory requirements.

2.3 Cross-domain trust: Cross-Domain Trust and Heterogeneous Trust Roots

Cross-domain trust addresses agentic interactions that span permitted administrative or security-domain boundaries. It supports the controlled exchange and evaluation of trust evidence while preserving each participating domain's trust anchors, policy authority, security controls, and regulatory obligations.

Heterogeneous Trust Anchors and Evidence

Trust anchors provide an important foundation for establishing trust in mobile networks. Existing hardware-backed and operator-controlled trust mechanisms, including UICC/eUICC-based credentials, secure elements, hardware roots of trust and PKI-based mechanisms, remain fundamental to mobile-network security. Agentic Networks do not remove the need for these mechanisms. Instead, increasingly diverse interactions among

Network AI Agents, AI-enabled network functions, third-party AI agents and external administrative domains may require MNOs to evaluate additional forms of trust evidence alongside established authentication, authorisation and security controls.

M-Trust therefore supports a heterogeneous trust foundation in which different trust anchors, credentials and verification mechanisms can contribute to trust evaluation according to the participating entities, administrative domains, required assurance level, risk, operator policy and applicable regulatory requirements. Hardware-backed trust anchors may provide strong protection for credentials and other security-critical assets, while appropriately assured software-based mechanisms may provide complementary protection in scenarios and at assurance levels where their use is permitted by operator policy. Such software-based mechanisms should not be presented as substitutes for hardware-backed protection where the latter is required. Cryptographic mechanisms, including digital signatures, integrity protection and other cryptographic proofs, provide means of verifying identity claims, data integrity, provenance or other trust evidence; they should not be considered independent alternatives to the underlying hardware or software trust foundations on which their credentials and execution environments may rely.

In interactions spanning multiple administrative domains, additional federated or distributed trust mechanisms may be considered where no single participating MNO controls the complete trust relationship. Their role is to support the exchange or validation of trust evidence across permitted domain boundaries, rather than to replace the established trust anchors, PKI, authentication mechanisms or security authority of the participating operators. Any such mechanism remains subject to the policy, security, administrative and regulatory requirements of the respective domains.

For example, interactions between a Network AI Agent and network functions within a single MNO domain may rely primarily on operator-controlled identity, credentials, authorisation and hardware-backed trust mechanisms where required. A third-party AI agent requesting access to exposed network capabilities may additionally be required to provide verifiable trust evidence acceptable to the operator. Where an interaction spans multiple operator or administrative domains, the participating domains may establish mechanisms to exchange or validate trust evidence while each retains authority over its own network capabilities and security policies. The applicable trust foundation is therefore determined by the MNO security architecture and the requirements of the interaction, rather than being autonomously selected or replaced by the AI agent itself.

Cross-Domain Trust Establishment

Agentic interactions in future mobile networks may cross multiple administrative and security domains, for example when a Third-party AI Agent interacts with exposed MNO network capabilities, or when Network AI Agents and network functions interact across operator or service-provider domains. Crossing such a boundary does not remove the need for each domain to apply its own trust, security, and authorisation controls. Rather, the challenge is to enable relevant trust and authorisation context or claims to be securely conveyed and evaluated across permitted domain boundaries without requiring each domain to recreate all trust context from the beginning.

Cross-domain trust establishment may therefore rely on three complementary capabilities:



Interoperable trust credentials and evidence



Cross-domain trust validation mechanisms



Trust and authorisation context continuity

In an MNO context, a Third-party AI Agent requesting access to an exposed network capability may present verifiable identity, trust, authorisation, or compliance evidence issued or endorsed by another trusted domain. Similarly, interactions between operator domains may require relevant trust and invocation context to be conveyed across the administrative boundary. The receiving MNO evaluates this evidence according to its own operator policies, security controls, risk assessment, authorisation requirements, and applicable regulatory obligations before granting access to network capabilities or resources.

Cross-domain mechanisms may preserve relevant trust and authorisation context as an interaction progresses across permitted domains, reducing unnecessary duplication of trust establishment while maintaining domain separation. Trust established in one domain is therefore not automatically inherited by another domain, nor does cross-domain trust imply globally transferable authorisation. Each participating MNO or administrative domain retains authority over which external trust evidence it accepts, how that evidence is evaluated, and what network capabilities, if any, are authorised.

2.4 Relationships with Zero Trust

M-Trust is not intended to replace Zero Trust or existing mobile-network security mechanisms. It builds on relevant Zero Trust principles, including the absence of implicit trust, least-privilege access, continuous verification, and policy-based authorisation, and considers how these principles may be complemented for increasingly autonomous and agentic interactions in future mobile networks.

The introduction of Network AI Agents, AI-enabled network functions, and third-party AI agents does not replace existing users, devices, applications, or network functions. Instead, it introduces additional interaction patterns in which agents may plan and execute actions, delegate tasks, invoke other agents, act on behalf of users or applications, and request access to network capabilities across changing execution contexts. These characteristics may require MNOs to evaluate not only identity, device and resource-access information, but also the authorised objective, delegation and invocation chain, requested network capabilities, observed execution behaviour, trust evidence, and associated risk.

M-Trust therefore complements Zero Trust approaches in three areas:

01

It extends continuous evaluation to consider whether an agent's requested and executed actions remain aligned with the authorised objective throughout a potentially multi-step invocation chain.

02

It supports adaptive authorisation and policy enforcement based on agent-specific context, trust evidence, observed behaviour and risk, while preserving least privilege and operator-defined policy constraints.

03

It supports the controlled exchange and evaluation of trust evidence across permitted administrative boundaries, while each participating MNO or administrative domain retains authority over its own trust decisions, network capabilities and security controls.

Cross-domain trust in M-Trust does not imply universal trust, the transfer of an authorisation decision from one domain to another, or reliance on a single trust authority. Trust evidence obtained from another domain may inform a local decision, but it remains subject to independent verification and evaluation under the receiving domain's policies, layered security controls, assurance requirements and applicable regulatory obligations.

M-Trust should therefore be understood as a proposed framework for augmenting existing identity, authentication, authorisation and policy mechanisms with additional trust context, evidence, evaluation and governance for autonomous interactions within, with, and across the permitted boundaries of future mobile networks.

3 High-Level Architecture

3.1 Security and Resilience of the M-Trust Framework

M-Trust is intended to augment, rather than replace or centralise, existing MNO authentication, authorisation, policy enforcement and domain-specific security controls. Its logical functions should therefore not be interpreted as requiring a single centralised trust authority or a single decision point upon which the security of the entire network depends. Trust evaluation and enforcement may be deployed according to network, service and administrative boundaries, while the relevant MNO security controls and policy authority remain authoritative within each domain.

The M-Trust functions and the trust information they process are themselves security-critical assets. Their implementation should therefore consider protection of trust evidence and policy information against unauthorised modification, appropriate authentication and authorisation between M-Trust functions, separation of duties and least privilege, secure interfaces and communications, auditability, and resilience against the compromise or unavailability of individual components. A compromise of one M-Trust function or trust domain should not automatically establish trust in, or authorise access to, other network domains. Implementations should also provide appropriate failure isolation and blast-radius containment, so that anomalous behaviour, failure or compromise of an individual agent or M-Trust function remains contained and does not cause cascading failures across functions or domains.

These principles allow M-Trust to provide additional trust context and evaluation for agentic interactions without creating a universal trust authority or weakening the layered security controls of the underlying mobile network [13].

3.2 Overview of the Architecture

M-Trust aims to establish an evolvable, verifiable, and interoperable trust framework for future intelligent communication networks. Unlike traditional trust models that focus primarily on identity verification and access control, M-Trust treats trust assessment as context-dependent and subject to continuous re-evaluation for relevant agentic interactions. Trust assessment outcomes are managed as computable, continuously evolving, and adaptive system attributes that can support intelligent security decision-making.

To achieve this objective, the framework is designed according to three fundamental architectural principles:

Layered Collaboration — Different trust functions operate independently while remaining tightly coordinated across the architecture.

Bounded Automation — Trust evaluation and policy adaptation may be automated within predefined operator governance, security, and authorisation boundaries.

Closed-Loop Feedback — Trust evaluation and policy enforcement continuously interact through monitoring and risk-based adaptation within operator-defined governance and security boundaries. Routine adjustments may be automated where permitted by operator policy, while higher-risk, exceptional, or policy-sensitive decisions may require additional verification or escalation to operator-controlled or human oversight. Human oversight should be risk-based: routine actions within predefined policy boundaries should not normally require synchronous human approval for each interaction and may instead be supervised through asynchronous monitoring, auditing and sampling, while human or operator approval is concentrated at key points such as policy-boundary changes and high-impact or exceptional operations, so that human involvement does not become a performance bottleneck for agentic interactions. Human intervention remains available where required to review, override, suspend, or approve security-relevant decisions.

Based on these principles, the M-Trust architecture is organised into three functional layers:

Trust Collection and Evaluation Layer

Data Trust Indicator Collection (DTIC) | Distributed Trust Evaluation (DTE)

Policy Decision Layer

Policy Agent (PA, Policy Decision and Validation Function)

Policy Enforcement Layer

Policy Enforcement Agent (PEA)

The three layers exchange trust information through defined interfaces. To support interoperable implementation across network functions, Network AI Agents and administrative domains, standardised interface definitions and trust-information exchange formats, together with a common taxonomy and semantics for trust evidence, may be required. Any exchange of trust information should be subject to applicable legal, regulatory and policy requirements, including requirements relating to data protection, data transfer and competition law. Together, these interactions support an end-to-end trust management lifecycle. DTIC is responsible for collecting and protecting trust-related data; DTE performs trust assessment and risk analysis; the Policy Decision Function determines and validates applicable policy and authorisation conditions according to trust and contextual information; and PEA executes security controls and enforcement actions based on policy decisions. Through this architecture, M-Trust supports continuous trust management under highly distributed, uncertain, and dynamically changing environments.

3.3 Trust Collection and Evaluation Layer: Data Foundation and Trust Perception

The Trust Collection and Evaluation Layer serves as the foundation of the entire M-Trust architecture. Its purpose is to collect trust-relevant evidence, establish trustworthy observations, and generate trust assessment results that support higher-layer policy decisions.

Data Trust Indicator Collection(DTIC)

DTIC functions as the trust-sensing component of the framework. Its primary responsibility is to collect trust-related information from heterogeneous sources, including:

- Agent behavioural records
- System operational status
- Policy execution logs
- Security incidents and anomaly reports

The reliability of M-Trust decisions depends on the quality and integrity of the trust evidence used in their evaluation. DTIC should therefore ensure that collected trust evidence is appropriately authenticated, integrity-protected, attributable to its source where required, and sufficiently fresh for the relevant trust decision. These requirements apply within both single-MNO and multi-domain deployments. Depending on the deployment environment and required assurance level, existing operator-controlled security mechanisms may be complemented by mechanisms such as hardware-backed protection, remote attestation, digital signatures, and cryptographic verification.

Where trust evidence is exchanged across administrative domains, DTIC should additionally support verification of its provenance and the authority under which it was produced. Evidence received from another domain should not be considered trusted solely because it is cryptographically verifiable; its acceptability remains subject to the receiving MNO's trust policy, security controls, and applicable authorisation requirements.

Trust evidence may contain sensitive information, including identity attributes, observed agent behaviour, and access or execution information. DTIC should therefore apply appropriate data protection, minimisation, access control, and privacy measures in accordance with operator policy and applicable regulatory requirements [15]. More computationally intensive privacy-enhancing techniques may be considered for specific use cases where their security and privacy benefits justify the associated operational overhead, but they are not a general requirement of the M-Trust architecture.

Distributed Trust Evaluation(DTE)

Distributed Trust Evaluation evaluates trust evidence relevant to an agentic interaction and produces a contextual trust and risk assessment for use by operator-controlled policy and authorisation functions. Depending on the interaction, this evidence may include identity and credential status, attestation or assurance information, the authorised objective, delegation and invocation chains, requested network capabilities, execution context, observed behaviour, and relevant risk indicators. DTE complements rather than replaces existing MNO trust anchors, authentication, authorisation, policy, and security controls.

"Distributed" refers only to the possible distribution of evidence collection or evaluation functions. MNOs may adopt centralised, hierarchical, distributed, or federated deployment according to architectural, assurance, latency, resilience, and administrative requirements. Distribution should not be assumed to improve security by itself.

Within a single MNO administrative domain, trust evaluation may rely primarily on operator-controlled trust anchors, credentials, protected data sources, and existing security and policy mechanisms. Distribution of evaluation functions may be used where justified by latency, scalability, availability, local-data processing, or failure-containment requirements, but should not be assumed to improve security by itself. The participating functions, interfaces, evidence flows, and evaluation logic must be protected against compromise, manipulation, unauthorised access, and inconsistent policy application.

Where trust evidence or assessment results are exchanged across permitted administrative boundaries, each receiving domain remains responsible for validating the source, provenance, freshness, scope, and relevance of that information under its own policy. Evidence or an assessment produced in one domain does not automatically establish trust or authorise access in another domain. Federated or distributed validation mechanisms may be considered for specific multi-domain use cases where the participants, governance arrangements, fault assumptions, decision scope, and operational benefits are clearly defined.

Artificial intelligence, statistical methods, and rule-based techniques may support trust evaluation, but the choice of evaluation method is implementation-specific. The method should be appropriate to the required assurance level, explainability, robustness, responsiveness, and operational constraints of the relevant MNO scenario. M-Trust does not mandate DLT, federated learning, decentralised consensus, or any particular machine-learning model as a component of DTE.

3.4 Policy Decision Layer: Policy Evaluation and Adaptation

The Policy Decision Layer translates trust and risk assessments into policy and authorisation decisions within operator-defined governance, security, and regulatory constraints. It is a logical function of the M-Trust architecture and should not be interpreted as a centralised authority replacing existing MNO policy, authentication, authorisation, or security-control functions. Its role is to use additional trust context provided by M-Trust to support policy decisions and adaptations while preserving MNO policy authority.

Two complementary capabilities are provided within this layer:

Context-Aware Policy Decision + **Policy Validation and Oversight**

Context-Aware Policy Decision

Context-Aware Policy Decision determines the applicable authorisation scope and security conditions for an agentic interaction within predefined operator policy. Relevant inputs may include the authorised objective, identity and credential status, trust and risk assessments, delegation or invocation chains, requested network capabilities and resources, execution context, observed behaviour, administrative domain, and applicable policy constraints.

These inputs may support dynamic adjustment of authorisation scope, verification requirements, monitoring conditions, or other security controls as the trust and risk context changes. Lower-risk actions may use streamlined controls where permitted by operator policy, while access to sensitive network capabilities, higher-risk operations, or cross-domain interactions may require enhanced verification, additional trust evidence, more restrictive authorisation conditions, or operator-controlled approval. Such adaptation does not permit an AI agent to grant itself additional authority or to modify the underlying MNO security policy outside predefined governance boundaries.

Rule-based, analytical, or AI-assisted techniques may support these decisions, but M-Trust does not prescribe a particular decision algorithm. Where AI techniques are used, their outputs remain subject to operator-defined policy constraints and the required security and assurance controls.

Policy Validation and Oversight

Policy validation should occur before a new or materially modified policy decision is enforced. The validation process checks that the proposed decision is consistent with the authorised objective, applicable operator policy, authorisation constraints, security requirements, and relevant regulatory obligations. A decision that does not satisfy these conditions should not be enforced and may instead be denied, restricted, subject to additional verification, or escalated for operator-controlled or human review where required.

After enforcement, runtime monitoring provides feedback on changes in trust, risk, execution behaviour, resource usage, or other relevant context. Where these changes affect the basis of the original authorisation decision, the Policy Decision Layer may re-evaluate the decision and, within predefined policy boundaries, strengthen verification, restrict or revoke authorisation, introduce additional conditions, or trigger escalation. Higher-risk or exceptional changes may require additional testing, operator-controlled approval, or human intervention before they are applied.

Policy decisions, validation results, authorisation changes, and relevant interventions should be auditable and protected against unauthorised modification in accordance with MNO security and compliance requirements. This closed-loop process enables M-Trust to adapt policy enforcement to changing trust and risk conditions without transferring policy authority from the MNO to the autonomous agent or the M-Trust framework itself.

3.5 Policy Enforcement Layer: Policy Enforcement and Response

The Policy Enforcement Layer applies authorisation decisions and security controls determined through the Policy Decision Layer at the relevant enforcement points. It complements rather than replaces existing MNO security and enforcement mechanisms. Enforcement follows secure-by-design, least-privilege and default-deny principles: access to network capabilities or resources is permitted only where explicitly authorised under applicable operator policy.

Policy Enforcement Agents (PEAs) represent logical enforcement functions that may be deployed at appropriate network or service enforcement points. A PEA acts on validated policy and authorisation decisions from the Policy Decision Layer and may use relevant local context when applying those decisions. PEAs do not independently expand authorisation scope or create security policy. Any automated enforcement action remains within operator-defined policy, security and governance boundaries.

Adaptive Policy Enforcement

M-Trust supports risk-based adaptation of enforcement controls where permitted by operator policy. Changes in trust, risk, execution behaviour, or context may lead the Policy Decision Layer to require additional verification, restrict the authorisation scope, suspend or revoke access, or apply additional security conditions. The PEA enforces the resulting decision at the appropriate enforcement point.

For example, where an agent's observed execution deviates materially from its authorised objective or relevant trust and risk conditions change, an existing authorisation may be restricted or suspended. Sensitive network capabilities may require stronger authentication, additional trust evidence, or operator-controlled approval before access is granted or continued. Where supported by the underlying network security architecture, existing containment controls, including micro-segmentation in relevant virtualised 5G Core deployment scenarios, may be applied in accordance with operator policy and the identified threat [14].

Automated responses should be proportionate to their potential operational impact. Routine actions may be automated within predefined policy boundaries, while actions that could significantly affect network availability, privileged administration, or critical network functions may require additional safeguards, escalation, or human/operator intervention. Recovery from a restrictive action is subject to re-evaluation and applicable operator policy; restoration of previous authorisation should not be assumed solely because a trust or risk indicator has returned to its earlier state.

Failure-handling design should distinguish between agent access to network capabilities and the services carried by the network itself. Where M-Trust functions become anomalous or unavailable, new or expanded requests by agents to access network capabilities should be denied by default (fail-closed). At the same time, existing network services, in particular connectivity carrying public communications, should remain available under established operator security mechanisms and policies (fail-safe), and should not be disrupted solely because an additional M-Trust trust function has failed.

Automated protective measures such as rate limiting or circuit breakers may be used to stop runaway agent invocation loops or denial-of-service-like traffic patterns. Their trigger thresholds should be set carefully and calibrated over time, and may take into account the agent's authorised objective, behaviour and trust signals rather than relying only on raw traffic metrics, to avoid false triggering that causes self-inflicted service disruption or locks operator administrators out of critical systems.

Enforcement Feedback and Re-evaluation

Enforcement outcomes may provide additional evidence for subsequent trust and risk evaluation, including whether the intended control was successfully applied, relevant changes in execution behaviour, operational impact, and other security-relevant observations. Such feedback may support re-evaluation of an existing authorisation decision and, where appropriate, adjustment of its scope or conditions.

Feedback may also inform offline improvement of evaluation models, policies, or operational baselines subject to appropriate validation, change control, and operator governance. It should not be interpreted as allowing the Policy Enforcement Layer or an AI agent to autonomously modify the underlying MNO security policy. The resulting closed loop is therefore:

Trust Evidence Collection → Trust and Risk Evaluation → Policy and Authorisation Decision → Pre-enforcement Validation → Enforcement → Monitoring and Feedback → Re-evaluation

This enables adaptive enforcement while preserving existing MNO security controls, layered defence [13], operator authority, and appropriate human oversight.

3.6 Cross-Layer Collaboration and Trust Closed Loop

M-Trust's multi-layer architecture, the Trust Collection and Evaluation Layer, the Policy Decision Layer, and the Policy Enforcement Layer, does not operate as a set of isolated modules. Instead, within the M-Trust framework, and through the collaborative interaction of information flows and control flows, these layers collectively implement the closed

loop defined above. Each layer performs its own distinct role while exchanging relevant information with the others, supporting closed-loop trust evaluation and policy enforcement within operator-defined governance, security, and authorisation boundaries.

When DTIC in the Trust Collection and Evaluation Layer receives trust evidence or an externally generated alert indicating anomalous behaviour or potentially manipulated data, it does not need to wait for the regular evaluation cycle — it can directly trigger a cross-layer alert signal. The anomalous data characteristics (such as forged-credential signature traces or the time-series pattern of the anomalous behaviour) are pushed to DTE in real time, triggering DTE to launch an emergency re-evaluation process that prioritises re-assessing the subject's trust state and risk level based on the anomalous data, and produces an updated trust and risk-assessment.

Based on this assessment, the Policy Decision Layer determines whether the applicable authorisation scope or security conditions should be adapted within operator-defined policy constraints. Any new or materially modified policy decision is subject to pre-enforcement validation against the authorised objective, applicable operator policy, authorisation constraints, and security requirements. Once validated, the Policy Enforcement Agent (PEA) applies the resulting enforcement action in accordance with operator policy, monitors relevant execution outcomes, and provides feedback for subsequent trust and risk evaluation. Such feedback may also inform offline improvement of evaluation models or policy configurations where appropriate, subject to validation, change control, and operator governance.

Through the interaction of information and control flows across the layers, M-Trust enables coordinated cross-layer operation and supports adaptation between changing trust and risk conditions and policy enforcement. This closed-loop process complements existing MNO security controls and enables trust and authorisation decisions to be re-evaluated as relevant evidence and context change, while preserving MNO policy authority and appropriate human or operator oversight.

4 Technology Considerations for M-Trust

M-Trust is a technology-neutral framework intended to augment, rather than replace, existing MNO authentication, authorisation, policy enforcement, trust anchors, and domain-specific security controls [10]. Its implementation may draw on established and emerging technologies for the protection of trust evidence, trust and risk evaluation, agent identity and authorisation, secure invocation, and controlled interaction across administrative domains. However, no single technology, deployment model, or trust-establishment mechanism is mandatory across all M-Trust scenarios.

Technology choices should reflect the participating entities, deployment location, administrative boundaries, required assurance level, operational risk, performance constraints, operator policy, and applicable regulatory requirements. Existing operator-controlled mechanisms, including UICC/eUICC-based credentials where applicable, PKI, authentication and authorisation mechanisms, network-function security, and layered security controls, remain fundamental. M-Trust provides additional trust context, evidence, evaluation, and governance for increasingly autonomous interactions within, with, and across the permitted boundaries of future mobile networks.

This chapter therefore describes implementation considerations rather than prescribing a fixed technology stack. Emerging technologies may support particular M-Trust requirements, but their applicability should be assessed against the operational, security, governance, and performance requirements of the relevant MNO scenario.

4.1 Relationship with Existing MNO Security Mechanisms

M-Trust should be implemented as an additional trust-evaluation and governance capability within the wider MNO security architecture. It should not establish an alternative identity system or universal trust authority, nor should it require existing access-security, network-function security, credential-management, or policy-control mechanisms to be replaced.

For interactions within a single MNO administrative domain, trust may primarily rely on operator-controlled identity, credentials, authentication, authorisation, policy, and security mechanisms. Hardware-backed trust anchors, secure elements, UICC/eUICC-based credentials, and established PKI mechanisms may contribute to the protection of identities, credentials, keys, workloads, and network functions where applicable. GSMA guidance already identifies

the SIM as a standardised hardware Root of Trust for relevant IoT device-to-cloud security scenarios [12]; M-Trust should therefore build on, rather than displace, such established mechanisms.

Additional M-Trust capabilities become relevant where agentic interactions introduce longer invocation chains, delegation between agents, changing execution contexts, access to exposed network capabilities, or interaction across multiple administrative domains. In these cases, existing security mechanisms may be complemented with additional evidence concerning the authorised objective, agent role, delegation scope, invocation chain, execution behaviour, platform state, and associated risk.

Centralised, hierarchical, distributed, and federated deployment models may all be appropriate under different conditions. Distribution should not be assumed to improve security by itself, as additional nodes, interfaces, credentials, and synchronisation arrangements may also expand the attack surface. The selected deployment model should support the required scale, latency, resilience, assurance, and failure-containment objectives while preserving the relevant MNO's policy and security authority.

4.2 Protection and Assurance of Trust Evidence

The reliability of M-Trust decisions depends on the quality and integrity of the trust evidence used in their evaluation. Trust evidence should therefore be appropriately authenticated, integrity-protected, attributable to its source where required, and sufficiently fresh and relevant for the decision being made.

Depending on the required assurance level and deployment environment, established operator-controlled mechanisms may be complemented by hardware-backed key protection, secure execution environments, remote attestation, digital signatures, protected logging, and other cryptographic verification mechanisms. These mechanisms provide different forms of assurance and should not be treated as interchangeable trust roots. For example, attestation may provide evidence about a platform or workload state, but the relying MNO still needs to determine whether the issuer, evidence, measurements, and associated claims are acceptable under its policy.

Trust evidence may contain sensitive information, including identity and credential attributes, agent behaviour, invocation history, execution context, and access records. Technical permission to exchange trust information does not itself imply that such information may legally be exchanged. Appropriate data minimisation, access control, retention, confidentiality, and privacy protections should therefore be applied in accordance with operator policy and applicable legal and regulatory requirements, including data protection, data transfer and competition law requirements, where applicable[15].

Where computationally intensive privacy-enhancing technologies are considered, their use should be justified by a clearly defined requirement and assessed against latency, scalability, complexity, and operational overhead. Such technologies are potential options for particular use cases rather than general requirements of the M-Trust architecture.

Trust evidence and policy information are themselves security-critical assets. The functions that collect, exchange, evaluate, and use this information should be protected through least privilege, separation of duties, appropriate information firewalls to restrict unauthorised or inappropriate information flows, secure interfaces and communications, auditability, resilience, and appropriate lifecycle management. Compromise of one M-Trust function or trust domain should not automatically establish trust in, or authorise access to, another network domain.

4.3 Trust Evaluation and Policy Decision Techniques

M-Trust may use rule-based, statistical, analytical, or AI-assisted techniques to evaluate trust evidence and support policy and authorisation decisions. The choice of technique is implementation-specific and should reflect the assurance, robustness, explainability, responsiveness, and operational constraints of the relevant MNO scenario.

The output of trust evaluation should inform, rather than replace, operator-controlled policy and authorisation. A trust assessment or risk indicator does not itself constitute authorisation. The Policy Decision Layer should evaluate the available trust and risk information together with the authorised objective, requested network capabilities, delegation and invocation chain, execution context, administrative domain, operator policy, and applicable security and regulatory constraints.

Where AI-assisted techniques are applied to trust evaluation or policy decisions, their outputs should remain bounded by operator-defined rules and governance. Suitable controls may include input and output validation, policy constraints, testing, monitoring, auditability, fallback behaviour, and human or operator oversight where required. AI-assisted techniques should not autonomously expand authorisation scope, redefine operator policy, or grant an agent additional authority.

New or materially modified policy and authorisation decisions should be validated before enforcement. Runtime monitoring should then determine whether the trust, risk, execution behaviour, or context supporting the original decision has materially changed. Where necessary, authorisation scope or conditions may be restricted, additional evidence or verification may be required, or access may be suspended or revoked. These mechanisms should be implemented consistently with the security, validation, auditability, and change-control principles described in Chapter 3.

4.4 Agent Identity, Delegation, and Secure Invocation

M-Trust should distinguish between mobile-network access authentication and the identity and authorisation of AI agents operating within, on behalf of, or in interaction with the network. Existing or 3GPP access-security mechanisms, and their future evolution, continue to protect network-access relationships. Agent identity, delegation, and invocation security represent additional considerations and should not be used to imply that mechanisms such as 5G-AKA, network-function security, or established PKI arrangements are obsolete.

Network AI Agents should be subject to an operator-controlled lifecycle covering authorisation to create or deploy agents, onboarding, identity and credential assignment, permitted functions, authorisation scope, software or model changes where relevant, monitoring, and revocation or decommissioning. An agent acting on behalf of a user, application, network function, or another agent should have a verifiable delegation relationship and a defined authorisation scope. The identity and authority of the delegated agent should not be inferred solely from the identity of the originating requester.

Third-party AI Agents requesting access to exposed network capabilities should be authenticated and authorised through operator-approved mechanisms. Access should follow secure-by-design, least-privilege, and default-deny principles: network capabilities are made available only where explicitly authorised under applicable MNO policy. Higher-risk requests may require additional trust evidence, stronger verification, more restrictive authorisation conditions, or operator-controlled approval.

Agent-to-agent and agent-to-function interactions may involve multi-step invocation chains. Security controls should therefore support the protection of request integrity, identification of the invoking and delegated entities, validation of the authorised objective and relevant authorisation scope, and auditability of security-relevant actions. Each receiving network function or administrative domain remains responsible for deciding whether the requested action is authorised.

Accountability is an important component of trust and authorisation governance in Agentic Networks. The GSMA Responsible AI Maturity Roadmap includes accountability, together with human agency and oversight and transparency, among the key principles of responsible AI [17]. In M-Trust, accountability means that significant agent actions relating to security, authorisation or network resources should be traceable to clearly identified responsible and authorising parties, such as the MNO deploying a Network AI Agent, the provider of a third-party AI agent, and the user, application or other party that initiated or authorised the task. To this end, the authorised objective, delegation and invocation chains, relevant trust evidence, policy decisions and execution outcomes should be recorded in a tamper-resistant and auditable manner, enabling the MNO to reconstruct key decision processes after the fact, clarify the responsibilities of the parties involved and, where necessary, take corrective action, roll back changes, revoke permissions or apply other remedies. When anomalies occur, these records should also enable the MNO to perform root cause analysis promptly, for example to determine whether an anomaly resulted from erroneous model reasoning, faulty input data, an incorrect tool or cross-domain capability invocation, or other factors. Agent autonomy should not create an accountability gap: an agent may act autonomously within its authorised boundaries, but the corresponding responsibility and governance obligations remain with clearly identified people or organisations.

Future agent communication protocols may provide standardised support for agent discovery, capability description, delegation, invocation, and evidence exchange. However, this area remains under development. Current 3GPP work identifies communication involving UE-based agents, Network AI Agents, Third-party AI Agents, and agents within the 6G Core as relevant areas for study, rather than as a completed normative protocol framework. M-Trust should therefore remain protocol-neutral and support coexistence with existing 5G systems, future 6G systems, and mixed-generation deployments.

4.5 Cross-Domain Trust Evidence Exchange

Cross-domain M-Trust interactions may involve MNOs, enterprise environments, cloud or edge providers, service providers, and third-party AI Agents. These interactions take place across permitted administrative boundaries and remain subject to the security, policy, authorisation, and regulatory requirements of each participating domain. Cross-domain trust should support the controlled exchange and verification of relevant trust evidence. Such evidence may include credentials, attestations, authorisation claims, provenance information, or other security-relevant assertions. Cryptographic verification may demonstrate that evidence was issued by a particular source and has not been modified, but it does not require the receiving MNO to accept the evidence or grant access.

Each receiving domain should independently evaluate:

- issuer identity, authority, provenance, freshness and integrity;
- scope, subject, objective, capability and assurance level;
- consistency with receiving-domain policy, authorisation, security and regulatory requirements.

Accordingly, subject to applicable legal and regulatory requirements, trust evidence may be exchanged across permitted boundaries, but trust decisions and authorisation authority remain domain-specific. An authorisation decision made by one domain should not automatically authorise access to capabilities or resources controlled by another domain. This preserves layered defence and prevents cross-domain cooperation from creating a universal trust authority or single point of security failure.

Federated or distributed validation mechanisms may be considered where multiple administrative domains need to exchange or validate evidence and where no single participant controls the complete relationship. Their use should be based on clearly defined participants, governance arrangements, fault assumptions, evidence semantics, decision scope, and operational benefits. They complement rather than replace the established trust anchors and security authority of each participating domain.

4.6 Emerging and Use-Case-Specific Technologies

A range of emerging technologies may contribute to future M-Trust implementations. Examples include confidential computing, advanced privacy-enhancing technologies, post-quantum cryptography, distributed or ledger-based evidence recording, federated evaluation, and AI-assisted security techniques. These technologies are not mandatory components of M-Trust and should not be assumed to be suitable for every MNO scenario.

Their applicability should be evaluated against the specific MNO use case, compatibility with existing security architecture, assurance benefit, performance and operational overhead, attack surface, governance requirements, maturity, interoperability and applicable regulation.

M-Trust therefore remains independent of any particular emerging technology. The framework defines the trust, authorisation, governance, and evidence requirements for agentic interactions; individual MNOs and participating domains determine the appropriate implementation technologies within their respective security architectures and authority boundaries.

5 Conclusion and Future Outlook

5.1 Conclusion

M-Trust is not intended to replace Zero Trust, existing MNO security mechanisms, or current approaches to autonomous network operation. It builds on established principles including explicit authorisation, least privilege, continuous verification, layered security and policy-based control, and considers the additional trust and authorisation requirements that may arise as mobile networks evolve towards increasingly autonomous and agentic operation.

Autonomous networks already use intent-driven operation, AI-assisted decision-making and closed-loop automation to improve network and service management. M-Trust addresses a complementary question: as autonomous agents decompose tasks, delegate actions, invoke network capabilities and interact across changing contexts, how can MNOs continue to determine whether the entities involved remain trusted and whether their actions remain within the scope of the authorised objective?

Through Multi-party trust, Agentic Intent-aware trust and Cross-domain trust, M-Trust provides a conceptual framework for considering trust evidence and evaluation, context-aware authorisation, delegation and invocation chains, adaptive policy enforcement, and controlled trust-evidence exchange across permitted administrative boundaries. These capabilities are intended to augment existing operator-controlled identity, authentication, authorisation, policy and security mechanisms. MNO policy, security controls, administrative authority and applicable regulatory requirements remain authoritative.

M-Trust is proposed as a reference framework and research direction, not as a standard specification, mandatory architecture or prescribed technology stack. The concepts presented in this white paper require further validation through representative MNO use cases, architectural analysis, prototyping, security assessment, regulatory requirements and relevant industry and standards activities.

5.2 Future Outlook

The integration of agentic capabilities into future mobile networks remains an evolving area of industry and standards research. Current 3GPP work is studying how AI and agent capabilities may be supported in 6G, including interactions involving Network AI Agents, Third-party AI Agents and AI-enabled network functions. The resulting architecture, protocols, security mechanisms and deployment models are still developing. M-Trust should therefore evolve alongside this work rather than assume a particular future network architecture.

Several areas merit further investigation.

01 M-Trust needs to be mapped more precisely to mobile-network architecture and existing security mechanisms. Future work should determine where trust evidence is generated and evaluated, how M-Trust functions interact with existing authentication, authorisation and policy controls, and how Network AI Agents and third-party agents are managed over their lifecycle. Particular attention should be given to compatibility with existing 5G deployments and the evolution towards future 6G capabilities, rather than assuming that existing security architectures are replaced.

02 Trust and authorisation governance for agentic interactions requires further study. Autonomous-network frameworks already support intent-driven and closed-loop operations; M-Trust should focus on the additional requirements introduced when agents decompose tasks, delegate authority, invoke other agents or network capabilities, and change their execution path while pursuing an authorised objective. Further work is needed to determine how authorisation scope, delegation, invocation chains and trust evidence can be represented and evaluated without requiring prediction of every possible future agent action. For example, future work could define a structured, machine-readable representation of intent and the corresponding authorisation scope, enabling policy enforcement to rely on explicitly defined constraints rather than on an AI model's subjective interpretation of intent. TM Forum already positions intent-driven decision-making and closed-loop management as capabilities of Level 4 autonomous networks, reinforcing the need for M-Trust to focus on the trust and authorisation governance dimension rather than duplicating network automation [11].

03 Cross-domain trust requires clearly defined administrative and security boundaries. Future work should examine how trust evidence can be exchanged and validated among operators, enterprises, cloud or edge providers and third-party AI services while allowing each participating domain to retain authority over its own network capabilities and authorisation decisions. Cross-domain trust should not assume universal trust propagation or automatic transfer of authorisation between domains.

04 The security and resilience of M-Trust itself require further validation. Trust evidence, evaluation functions, policy decisions and interfaces may become security-critical assets. Future work should therefore evaluate attack surfaces, compromise containment, availability, separation of duties, auditability and the consequences of incorrect or manipulated trust assessments. Centralised, hierarchical, distributed and federated deployment models should be evaluated according to the requirements of specific MNO scenarios rather than assuming that one model is inherently more secure.

05 The balance among security assurance, responsiveness and operational complexity requires further validation. Trust evaluation and authorisation controls should remain appropriate to the network capabilities they protect, while emerging implementation technologies should be assessed against performance overhead, maturity, interoperability and compatibility with existing MNO security mechanisms.

As Agentic Network concepts and related standards evolve, M-Trust can provide a structured basis for industry discussion on how trust and authorisation governance should evolve with increasing network autonomy. Future work should progressively validate the framework against concrete MNO scenarios and emerging standards, refining its functional boundaries and identifying where new mechanisms are required and where established mobile-network security mechanisms remain sufficient.

6 Abbreviations

Abbreviation	Full Form
3GPP	3rd Generation Partnership Project
5G-AKA	5G Authentication and Key Agreement
AI	Artificial Intelligence
API	Application Programming Interface
DLT	Distributed Ledger Technology
DTE	Distributed Trust Evaluation
DTIC	Data Trust Indicator Collection
eSIM	Embedded SIM
eUICC	Embedded Universal Integrated Circuit Card
GSMA	GSM Association
IMT	International Mobile Telecommunications
IoT	Internet of Things
ITU	International Telecommunication Union
ITU-R	ITU Radiocommunication Sector
M-Trust	MAIC-Trust Framework
MNO	Mobile Network Operator
NIST	National Institute of Standards and Technology
PA	Policy Agent (Policy Decision and Validation Function)
PEA	Policy Enforcement Agent
PKI	Public Key Infrastructure
SIM	Subscriber Identity Module
SLA	Service Level Agreement
SP	Special Publication (NIST)
TR	Technical Report
TS	Technical Specification
UE	User Equipment
UICC	Universal Integrated Circuit Card

7 Terms and Definitions

Term	Definition
AI Agent	An intelligent entity that can perceive or obtain relevant environmental and contextual information and, through reasoning, planning and decision-making, execute tasks autonomously or in collaboration with other agents to achieve specific goals (see 3GPP TR 22.870 [18]).
Agentic Network	A future network in which network functions, Network AI Agents and third-party agents collaborate autonomously, achieving authorised objectives through goal decomposition, delegation and capability invocation.
Autonomous Networks	Networks that apply intent-driven operation, AI-assisted decision-making and closed-loop automation to improve network and service management, as positioned by TM Forum.
Zero Trust	A security model that grants no implicit trust based on network location and builds on identity-centric security, least privilege and continuous monitoring (NIST SP 800-207).
M-Trust (MAIC-Trust Framework)	The trust framework proposed in this white paper for Agentic Networks. It extends Zero Trust with three dimensions: Multi-party trust, Agentic Intent-aware trust and Cross-domain trust. It is positioned as a reference framework and research direction, not a standard or mandatory implementation architecture.
Multi-party trust	Trust evaluation in which multiple authorised parties contribute independently verifiable, purpose-bound and time-limited trust evidence, combined with context-appropriate trust establishment. External evidence informs, but does not replace, the MNO's local policy decision.
Agentic Intent-aware trust	Trust in AI agents built on identity verification and intent-aware authorisation, with policies adapted within operator-defined boundaries according to risk and context.
Cross-domain trust	Controlled exchange and evaluation of trust evidence across permitted administrative or security-domain boundaries, while each participating domain retains its own trust anchors, policy authority, security controls and regulatory obligations.
Intent-aware Authorisation	Authorisation that, beyond verifying identity, assesses whether a requested action remains consistent with the authorised objective, the delegation and invocation chain, the execution context and applicable policy. Actions beyond the pre-authorised scope trigger a new authorisation decision.
Authorised Objective	The task goal an agent is permitted to pursue, together with its permission boundaries.
Delegation and Invocation Chain	The sequence through which a task is delegated from one agent to another, and through which agents call other agents, tools or network capabilities. Both form part of the context evaluated in intent-aware authorisation.
Trust Evidence	Verifiable information used to support trust evaluation, such as identity and subscription data, observed network behaviour, workload or platform attestation, software integrity and delegated permissions.
Trust Anchor	A foundational element from which trust is established and against which credentials and trust evidence are verified, such as UICC/eUICC-based credentials, secure elements, hardware roots of trust and PKI.
Administrative Domain	The scope of networks or systems over which a single entity, such as an MNO or service provider, holds policy, security and management authority.
Network AI Agent	An AI agent deployed within, or acting on behalf of, an MNO network and governed by an operator-controlled lifecycle.
Third-party AI Agent	An AI agent provided by a party other than the MNO that interacts with the mobile network, for example by accessing exposed network capabilities through APIs.

8 References

[1]. 3GPP, TR 23.801-01, “Study on Architecture for 6G System; Stage 2,” Release 20, work in progress.

[2]. M. Chen, Z. Yang, T. Qiu, J. Zhang, M. Shikh-Bahaei, and T. Ohtsuki, “AI-Native Networking: A Vision and Survey,” *IEEE Communications Surveys & Tutorials*, vol. 26, no. 1, pp. 1–34, 2024.

[3]. ITU-R, Recommendation M.2160-0, “Framework and overall objectives of the future development of IMT for 2030 and beyond,” November 2023.

[4]. Z. Xiao, X. Song, W. Ni, X. Tao, and G. Li, “Integrated Sensing and Communication: A Comprehensive Survey,” *IEEE Communications Surveys & Tutorials*, vol. 25, no. 2, pp. 790–823, 2023.

[5]. S. Rose, O. Borchert, S. Mitchell, and S. Connelly, “Zero Trust Architecture,” NIST Special Publication 800-207, August 2020.

[6]. 3GPP, TR 33.894 V18.0.0, “Study on applicability of the Zero Trust Security principles in mobile networks,” Release 18, September 2023.

[7]. 3GPP, TR 33.794 V19.1.0, “Study on enablers for Zero Trust Security,” Release 19, March 2025.

[8]. Y. Wang, X. Kang, T. Li, H. Wang, C.-C. Kang, and Z. Lei, “SIX Trust for 6G: Towards a Secure and Trustworthy 6G Network,” *IEEE Access*, vol. 11, pp. 107657–107668, 2023.

[9]. Y. Wang, X. Kang, and T. Li, “Towards an Open Trust Establishment Infrastructure for Future Network: Motivations, Models and Technologies,” *IEEE Access*, vol. 12, pp. 111196–111205, 2024.

[10]. 3GPP, TS 33.501 V18.8.0, “Security Architecture and Procedures for 5G System,” Release 18, January 2025.

[11]. TM Forum, “Autonomous Networks: Level 4 Industry Blueprint,” June 2024.

[12]. GSMA, “IoT SAFE: Robust IoT Security at Scale: The Why, What and How of Securing IoT Applications and Data,” June 2021.

[13]. GSMA, “GSMA Mobile Telecommunications Security Landscape 2025,” February 2025.

[14]. GSMA, FS.61 Version 1.0, “Micro-Segmentation in 5G Core Network Resource Pool Guidelines,” April 2025.

[15]. GSMA, “Mobile Privacy Principles: Promoting Consumer Privacy in the Mobile Ecosystem,” February 2016.

[16]. GSMA, “GSMA Mobile Telecommunications Security Landscape 2026,” February 2026.

[17]. GSMA, “The GSMA Responsible AI Maturity Roadmap,” September 2024.

[18]. 3GPP, TR 22.870 V20.0.0, “Study on 6G Use Cases and Service Requirements,” Release 20, March 2026.

[19]. GSMA, “The Mobile AI whitepaper,” March 2026.

Contributing Organisations:

Huawei Technologies Co., Ltd.
China Telecommunications Corporation
China United Network Communications Co., Ltd.
China Mobile Communications Group Co., Ltd.
China Telecom Research Institute
Beijing University of Posts and Telecommunications
State Information Center
China Mobile Research Institute
China Mobile Group Design Institute Co., Ltd.
NSFOCUS Technologies, Inc.

Supporting Organisations:

China Radio and Television Mobile Network Co., Ltd.
China Mobile (Hangzhou) Information Technology Co., Ltd.

Drafting Team:

Li Tieyan, Liu Hong, Wei Kai, Shi Lin, Gu Hongxun, Wu Shanshan, Chang Liyuan, Zheng Zhi, Tao Ye, Wang Xiaoqing, Yang Kai, Wang Yong, Jing Jing, Guo Zhihong, Cheng Chonghao, Kang Xin, Zhang Yunfeng, Fan Xiaoji, Shao Meng, Zhang Chen, Zhang Xi, Zhang Kun, Liu Wenmao, Xu Mengjun, Du Gang, Li Yan, Wang Yue, Wu Chao



GSMA
1 Angel Lane
London
EC4R 3AB
United Kingdom

Tel: +44(0)20 7356 0600
Fax: +44(0)20 7356 0601